

RESPONSE

Claims 1-19 were pending in the Application. Claims 1, 7, 14, and 19 are amended, new claims 20-26 are added, and claims 5-6 and 16-18 are canceled by the present Amendment. Upon entry of the present Amendment, claims 1-4, 7-15, and 19-26 are pending and presented for consideration. Applicants respectfully submit that no new matter is introduced by the present Amendment.

Applicants note that a copy of the PTO Form 1449 filed on August 28, 2003, as initialed by the Examiner, was not attached to the above-referenced Office Action mailed on January 2, 2004.

The Specification has been amended to correct typographical errors that one skilled in the art would have recognized as obvious errors, and for which one skilled in the art would have recognized the appropriate correction.

Claims 1, 7, 14, and 19 have been amended to recite an "I/O node." Support for the amendments may be found, for example, in originally filed claims 6 and 16-18; and in the Specification at page 23, lines 22, through page 24, line 7.

New claims 20-26 have been added. Support for the new claims may be found, for example, in originally filed claims 14-19; in Figures 6 and 7; and in the Specification at page 27, lines 21, through page 32, line 8.

Rejections Under 35 U.S.C. §102

Claims 1, 2, 3, 5, 8, 11, 12, 14, 16, 17, and 19 were rejected under 35 U.S.C. §102 as being anticipated by U.S. Patent No. 6,643,763 to Starke et al. (hereinafter "Starke"). Applicants respectfully traverse the foregoing rejections.

Starke

Starke teaches a "[m]ethod, system and program storage device . . . for implementing a register pipe between processing engines of a multiprocessor computing system." Abstract

(emphasis added). “Data is transferred between the first processing engine and the second processing engine through the register pipe without passing through the memory.” Abstract (emphasis added). “A control mechanism is provided within each processing engine to dynamically enable or disable the register pipe coupling any two processing engines of the multiprocessor computer system.” Abstract (emphasis added).

The “computing system 100 [of Starke] may operate in any of single instruction stream single data stream (SISD), single instruction stream multiple data stream (SIMD) or multiple instruction stream multiple data stream (MIMD) modes according to operation of and flow of information between processors 110.” Col. 6, lines 58-63 (emphasis added).

Figure 8 depicts Starke’s register pipes connecting four CPUs. Figure 8, and col. 12, lines 32-48. “Each enabled register pipe is identified by the pair of the CPU IDs that the pipe connects.” Col. 12, lines 38-39. Figure 9 depicts Starke’s connecting fabric connecting four CPUs. Figure 9, and col. 13, lines 20-28. “Shared connections 24 couple connecting fabric 22 to each processing engine 12.” Col. 13, lines 23-24. A feature common to Starke’s FIG. 8 and FIG. 9 is that “within each processing engine registers are selectively mapped to create the desired pipe register.” Col. 13, lines 30-32. As Starke notes, “[f]or low levels of scaling, the fully connected network of FIG. 8 allows the fastest and simplest operation, while for higher levels, the embodiment of FIG. 9 may be more practical. Col. 13, lines 42-45.

In summary, Starke discloses different types of connections between processing engines of a multiprocessor computer system “which allow direct transfer of information from a first processing engine to a second processing engine without requiring that the information pass through main memory of the computer system.” Col. 14, line 64, through col. 15, line 3 (emphasis added).

Starke Does Not Anticipate Independent Claim 1 as Currently Amended

With respect to original independent claim 1, the Office Action suggests that Starke discloses a “processing system that is comprised of a plurality of processing elements executing

identical instruction streams . . . a target in communication with one of the processing elements . . . [and] a connecting fabric . . . communicating the transactions between the processing elements and the target.” (Page 3.) The Office Action also notes that a connecting fabric “is interpreted as a switching fabric.” (Page 3.)

Claim 1 is a fault-tolerant data processing apparatus that, as currently amended, recites:

a plurality of data processing elements executing substantially identical instruction streams substantially simultaneously;
an I/O node in communication with at least one of the plurality of data processing elements; and
a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node.

Starke does not teach or suggest at least “an I/O node in communication with at least one of the plurality of data processing elements; and a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node” as recited in currently amended claim 1. Applicants therefore respectfully request that the rejection of claim 1 under 35 U.S.C. §102 be reconsidered and withdrawn.

Starke Does Not Anticipate Independent Claim 14 as Currently Amended

With respect to original independent claim 14, the Office Action suggests that Starke discloses a “a plurality of processing elements generating replicated execution to target addresses . . . [and] the use of a connecting fabric for communicating the transactions.” (Page 4.) The Office Action also suggests that replicated execution to target addresses “is interpreted as a identical transactions.” (Page 4.)

Claim 14 is a method for fault tolerant digital data processing that, as currently amended, recites:

- (a) generating by a plurality of data processing elements identical transactions each having an I/O node address; and
- (b) communicating the identical transactions asynchronously on a switching fabric to the I/O node identified by the I/O node address.

Starke does not teach or suggest at least “generating by a plurality of data processing elements identical transactions each having an I/O node address;” or “communicating the identical transactions asynchronously on a switching fabric to the I/O node identified by the I/O node address” as recited in currently amended claim 14. Applicants therefore respectfully request that the rejection of claim 1 under 35 U.S.C. §102 be reconsidered and withdrawn.

Starke Does Not Anticipate Claims That Depend From Claim 1 or Claim 14

Claims 2, 3, 5, 8, 11, and 12 depend from claim 1 and include all of the limitations of claim 1. Similarly, claims 16, 17, and 19 depend from claim 14 and include all of the limitations of claim 14. Since, as explained above, Starke does not teach or disclose all of the elements of claim 1 or claim 14 as currently amended, Starke cannot teach or disclose all of the elements of the claims that depend from those claims.

Therefore, in light of the foregoing reasons and amendments to the claims, Applicants respectfully request that the rejections under 35 U.S.C. §102 be reconsidered and withdrawn.

Rejections Under 35 U.S.C. §103

Claims 4, 6, 7, 13, and 18 were rejected under 35 U.S.C. §103 as being unpatentably obvious over Starke in view of U.S. Patent No. 6,643,764 to Thorson et al. (hereinafter “Thorson”). Claims 9 and 15 were rejected under 35 U.S.C. §103 as being unpatentably obvious over Starke in view of U.S. Patent No. 5,903,717 to Wardrop (hereinafter “Wardrop”). Claim 10 were rejected under 35 U.S.C. §103 as being unpatentably obvious over Starke in view of U.S. Patent No. 6,622,193 to Avery (hereinafter “Avery”). Applicants respectfully traverse the rejections as applied to the original and the amended claims, and submit that none of the references, alone or in combination, teach or suggest the claimed invention.

Thorson

Thorson is directed to “routing messages on multiple links in multiprocessor computer systems.” Col. 1, lines 9-10. Thorson addresses the need for improved infrastructure in “[m]ultiprocessor computer systems having up to hundreds or thousands of processing element nodes . . . typically referred to as massively parallel processing (MPP) systems.” Col. 1, lines 51-52 and lines 19-22.

In particular, Thorson teaches:

[A] multiprocessor computer system having a plurality of processing element nodes and an interconnect network interconnecting the plurality of processing element nodes. An interface circuit is associated with each one of the plurality of processing element nodes. The interface circuit has a lookup table having n-number of routing entries for a given destination node. Each one of the n-number of routing entries associated with a different class of traffic.

Col. 1, lines 55-63 (emphasis added). Thorson’s multiprocessor system “provides redundant paths.” Col. 5, lines 66-67. “The network traffic is routed according to class.” Abstract.

Fig. 3 depicts “one or more processing element nodes 302, 304 and an interconnect network 306 interconnecting the plurality of processing element nodes.” Figure 3 and col. 3, lines 16-19 (emphasis added). “Each one of the processing element nodes [in Fig. 3] has an interface circuit” that is distinct from the interconnect network. Col. 3, lines 19-20. Although the interface circuit in each of the processing element nodes is involved in routing I/O associated with the node, Thorson does not teach or suggest that the interconnect network communicates transactions between a processing element node and an I/O node. Figures 1 and 3, along with the related written description (emphasis added).

In summary, Thorson suggests routing traffic by class between processing element nodes in a multiprocessor computer system. Thorson does not teach or suggest any of the elements recited in claim 1 as currently amended, namely: “a plurality of data processing elements executing substantially identical instruction streams substantially simultaneously;” “an I/O node in communication with at least one of the plurality of data processing elements;” or “a switching

fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node.” Similarly, Thorson does not teach or suggest any of the elements of claim 14 as currently amended, namely: “generating by a plurality of data processing elements identical transactions each having an I/O node address;” or “communicating the identical transactions asynchronously on a switching fabric to the I/O node identified by the I/O node address.” Accordingly, neither Starke nor Thorson, alone or in combination, discloses all of the elements of claim 1, claim 14, or claims 4, 6, 7, 13, and 18, which depend from claim 1 or claim 14.

Wardrop

Wardrop generally discloses fault tolerant computing methods and apparatus that use redundant voting at the hardware clock level. Col. 5, lines 2-5. Wardrop specifically discloses a system based on a computer that contains three or preferably four CPUs. Col. 5, lines 19-24. In Wardrop’s system:

The output from each CPU is compared for agreement with the output from all other CPUs. Agreement of a majority of CPU output signals supplied to the voter results in a voted output signal which has the value of the majority. A CPU output signal which does not agree with the majority is detected by the voter, producing an error signal. The error signal is sent to the memory controller which reacts several ways:

1. The majority voted signal is used by the agreeing CPUs to continue CPU processing operations without interruption;
2. The disagreeing CPU is disabled from further participation in voting;
3. A system management interrupt (SMI) is generated to the other CPUs; and
4. At a later time, software initiates a re-synchronization process that recovers the disabled CPU.

In the event of failure of a computer, a spare, error-free computer is substituted.

Col. 5, line 53 through col. 6, line 4.

Wardrop teaches a system that substitutes a spare, error-free computer for the on-line CPUs when two CPUs are on-line and their output streams do not agree. Wardrop does not

disclose a fault tolerant computer system that includes the voter delay buffer disclosed in Figures 3 and 4 of the present application.

Accordingly, Wardrop states:

FIG. 5 depicts a state diagram 100 for one of two redundant computers 30 of the instant invention which has four CPUs 32. It shows system reaction to a fault detection and correction of the fault. When one CPU 32 of this embodiment is upset or suffers a failure, the computer 30 can still be operated with the remaining three CPUs 32. If a second CPU 32 is temporarily upset, the computer 30 will continue to function with two CPUs 32 in agreement, and all of the CPUs 32 will be resynchronized to bring all four CPUs 32 into agreement. *The logic of this preferred embodiment operates to substitute a spare computer 30 for the dysfunctional computer 30 in the event of permanent failure of two CPUs and the upset of a third.*

Column 10, lines 22-34 (emphasis added). Additionally Wardrop states:

While the computer is in “two CPUs agree state” 126, should a disagreement be detected among the two voting CPUs 32, the computer 30 will transition 134 to a “failed computer state” 136.

Upon reaching the “failed computer state” 136, the computer 31 indicates a failed condition to the reconfiguration unit and a transition 138 is made to “switch to an error-free redundant computer state” 140. The switch being made, the computer 30 is reset (normally all registers go to zero) and the system 10 is rebooted with the error-free, redundant computer 30. The substitute computer 30 then transitions 148 to a “four CPUs agree” state 102.

Column 11, lines 18-30. Wardrop further states:

If no two operating CPUs 32 agree 204, the system is considered to have failed 206. In such event, a system restart with a substituted spare computer 30, will be attempted. All software will be restarted, but a major disruption of computer functions will occur.

Column 13, lines 16-20 and FIG. 7.

In summary, Wardrop suggests a fault tolerant computing system that relies on the agreement of two or more CPU output streams to detect and to correct single event upsets (SEU) and other random failures. Wardrop does not teach or suggest the “plurality of voter delay buffers” recited in claim 9 or “a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node” as recited in

currently amended claim 1. Similarly, Thorson does not teach or suggest at least “communicating the identical transactions asynchronously on a switching fabric to the I/O node identified by the I/O node address” as recited currently amended claim 14 or “transmitting by the voting unit a single transaction asynchronously on a switching fabric” as recited in claim 15. Accordingly, neither Starke nor Wardrop, alone or in combination, discloses all of the elements of claim 1, claim 14, or claims 9 and 15, which depend from claim 1 or claim 14. .

Avery

The Office Action suggests that Avery discloses using a plurality of DMA in communication with a switching fabric. Page 7.

Avery discloses an infiniband system 300 that “has a memory interconnect bus 304 that connects CPU 302 and memory 306. Figure 3 and col. 5, lines 62-63. “The memory interconnect bus 304 is, in turn, connected to a host channel adapter (HCA) 308 that includes its own CPU 309 and memory 311.” Col. 5, lines 63-65.

In the system disclosed by Avery:

The HCA 308 is connected to a switch fabric 318 for both sending and receiving data as indicated schematically by arrows 314 and 316. The HCA 308 can be connected to any number of peripheral buses via the switch fabric 318. In particular, the HCA 308 can be connected to various PCI peripherals, of which two, 332 and 334 are shown, via a PCI bus 330, by means of a target channel adapter (TCA) 324. In this case, the TCA 324 is an Infiniband to PCI bridge (IB-PCI bridge) and can both send and receive data as indicated schematically by arrows 320 and 322. The TCA 324 also includes a CPU 325 and a memory 327. Other TCAs (not shown) may also be present.

Clients of both the HCA 308 and the TCS 324 can control data transfer by creating a facility called a work queue.

Col. 6, lines 7-19.

Avery shows “an overall view of a DMA data transfer process implemented in a message-passing, queue-oriented system such as that shown in FIG. 3.” Figure 5 and col. 8, lines 16-52.

Avery does not teach or suggest any of the elements recited in claim 1 as currently amended, namely: “a plurality of data processing elements executing substantially identical instruction streams substantially simultaneously;” “an I/O node in communication with at least one of the plurality of data processing elements;” or “a switching fabric communicating transactions asynchronously between at least one of the plurality of data processing elements and the I/O node.” Accordingly, neither Starke nor Avery, alone or in combination, discloses all of the elements of claim 1, or claim 10, which depends from claim 1.

Therefore, in light of the foregoing reasons and amendments to the claims, Applicants respectfully request that the rejections under 35 U.S.C. §103 be reconsidered and withdrawn.

SUMMARY

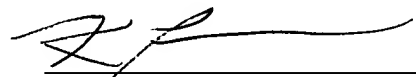
Claims 1-19 were pending in the Application. Claims 1, 7, 14, and 19 are amended, new claims 20-26 are added, and claims 5-6 and 16-18 are canceled by the present Amendment. Upon entry of the present Amendment, claims 1-4, 7-15, and 19-26 are pending and presented for reconsideration. Applicants respectfully submit that no new matter is introduced by the present Amendment.

Applicants request that the Examiner reconsider the application and claims in light of the foregoing Amendment and Response, and respectfully submit that the claims, as amended, are in condition for allowance. If, in the Examiner's opinion, a telephonic interview would expedite the favorable prosecution of the present application, the undersigned agent/attorney would welcome the opportunity to discuss any outstanding issues, and to work with the Examiner toward placing the application in condition for allowance.

Applicants submit a \$308 check, of which \$110 is designated to cover the one-month extension fee, \$18 is designated to cover the new claims added by the present Amendment, and \$180 is designated to cover the supplement IDS. Applicants believe that no additional fees are necessitated by the present Amendment. However, in the event that any additional fees are due, the Commissioner is hereby authorized to charge any such fees to Attorney's Deposit Account No. 20-0531.

Respectfully submitted,

Date: May 3, 2004
Reg. No.: 47,577
Tel. No. (617) 248-7593
Fax No. (617) 248-7100


Kia L. Freeman
Attorney for Applicants
Testa, Hurwitz, & Thibault, LLP
High Street Tower
125 High Street
Boston, MA 02110

3038480